Last name, first name: _____

Student number: _____

263-0007-00L: Advanced Systems Lab
ETH Computer Science, Spring 2021
Midterm Exam
Wednesday, April 21, 2021

Instructions

- Write your full name and student number on the front.
- Make sure that your exam is not missing any sheets.
- No extra sheets are allowed.
- The exam has a maximum score of 100 points.
- No books, notes, calculators, laptops, cell phones, or other electronic devices are allowed.

Problem 1 (18 = 2+2+4+6+4)	
Problem 2 (12 = 4+4+4)	
Problem 3 (20 = 2+2+6+6+4)	
Problem 4 (20 = 2+2+2+4+4+3+3)	
Problem 5 (18 = 7+5+6)	
Problem 6 (12 = 2+2+2+2+2+2)	
Total (100)	

Problem 1: Bounds (18 = 2+2+4+6+4)

Consider the following two functions:

```
1  void f1 (double *x, double *y, double *z, int n){
2    double t;
3    for(int i=0; i < n; i++){
4      double a = x[i];
5      double b = y[i];
6      double c = z[i];
7      z[i] = (5.0*a*a + 4.0*b*b) + 3.0*c;
8    }
9  }
10
11 void f2 (double *x, double *y, double *z, int n){
12   double t;
13   for(int i=0; i < n; i++){
14     double a = x[i];
15     double b = y[i];
16     double c = z[i];
17     z[i] = (a*a)/4.0 + b*b - c*c;
18   }
19 }
```

Assume that the above code is executed on a machine with the following relevant latency, gap (inverse throughput), and port information:

Instruction	Latency [cycles]	Gap (inverse throughput) [cycles/instruction]	Port
`add/subtract`	2	0.5	0/1
`mult`	3	1	1
`div`	5	4	2

The processor does **not** support vector instructions. Further assume that:

1. You can ignore the latency and throughput of loads and stores, i.e., assume they have zero latency and infinite throughput.

2. The compiler does not apply any algebraic transformation: the operations are mapped to assembly instructions as shown.

3. Ignore integer operations.

4. A division counts as one floating point operation.

Show enough detail with each answer so we understand your reasoning.

1. Determine the maximum theoretical floating point peak performance in flop/cycle of the machine under consideration.

2. Determine the flop count $W(n)$ of `f1` and `f2`.

3. Determine for each function: a lower bound (as tight as possible) for the runtime (in cycles) and an associated upper bound for the performance of `f1` and `f2` based on the instruction mix, ignoring dependencies between instructions (i.e., don't consider latencies and assume full throughput).

   `f1`:

   `f2`:

4. Suggest an improvement for each function, i.e. algebraic transformations, that produces the same result in real arithmetic but results in smaller lower bounds for their runtime. Again, only consider the instruction mix, ignoring dependencies. State also these new lower bounds for the runtime of `f1` and `f2`.

   `f1:`

   `f2:`

5. Estimate for the original function `f1` (i.e. not the one with improvement done in task 4) a lower bound for its runtime (as tight as possible) of one loop iteration taking latency, throughput and dependency information into account. Draw the corresponding DAG of one iteration of function `f1`.

Problem 2: Operational Intensity (12 = 4+4+4)

Consider the computation $z = Ax+Cy$ where x, y, z are column vectors of doubles of length n and A, C are $n \times n$ matrices of doubles. No temporary arrays are used in these computations. We assume a write-back/write-allocate cache with blocks of size 32 bytes and a cold cache at the start of the computation. `sizeof(double) = 8`. In the derivations you can omit lower order terms (writing $\approx$ instead of $=$). Show your work.

1. Determine a hard upper bound for the operational intensity $I(n)$ [flops/byte] of the computation considering only compulsory misses. Consider both reads and writes. **Note**: The upper bound should hold for all cases, independently of the memory layout of the operands and the implementation.

2. You are given a machine which has 2 ports. Each port can execute one add or one multiplication per cycle (no FMA, and no vector instructions). For what values of the memory bandwith β (in bytes/cycle) is the computation guaranteed to be memory bound?

3. Compute now a lower bound for the operational intensity assuming that every memory access leads to a misse. Exclude read/write acesses to z in the analysis.

Problem 3: Cache Mechanics (20 = 2+2+6+6+4)

You are given a cache with 4 sets and LRU replacement policy. Its block size is 16 bytes, and the capacity is 128 bytes. Consider the following code. `sizeof(double) = 8`.

```
1 double a[40];
2 double x = 0;
3 for (int i = 0; i < 20; ++i) {
4   double lhs = a[2*i];
5   double rhs = a[f(i)];
6   x += lhs * rhs;
7 }
```

Assume that array `a` starts at the memory address `0`. Variables `x`, `i`, `lhs` and `rhs` are stored in registers. Memory accesses happen in exactly the order that they appear. Answer the following. Show your work. Hint: It helps to draw the cache.

1. How many doubles fit into this cache?

2. What is the associativity of this cache?

3. For each of the following definitions of `f(i)` do the following two things: i) determine the miss rate; ii) draw the state of the cache at the end of the computation. Show your work.

   (a) `f(i) = (2*i) % 8`:

(b) `f(i) = (2*i + 4) % 40`:

4. Give an example of `f(i)` that yields the maximal miss rate. All memory accesses have to be valid (no index out of bounds accesses). Show your reasoning.

Problem 4: Roofline (20 = 2+2+2+4+4+3+3)

Assume a computer with the following features:

- A CPU that supports single precision floating point operations. The relevant port information is as follows:
  Port 1: fma, mul, add.
  Port 2: fma, mul.
- Each of these operations has a throughput of 1 per port and a latency of 4 cycles.
- It does not support any SIMD operations.
- A write-back/write-allocate cache of size 2 MiB with cache block size $B = 64$ bytes. The cache is initially cold.
- The read (memory) bandwidth is 8 floats per cycle. `sizeof(float) = 4`.

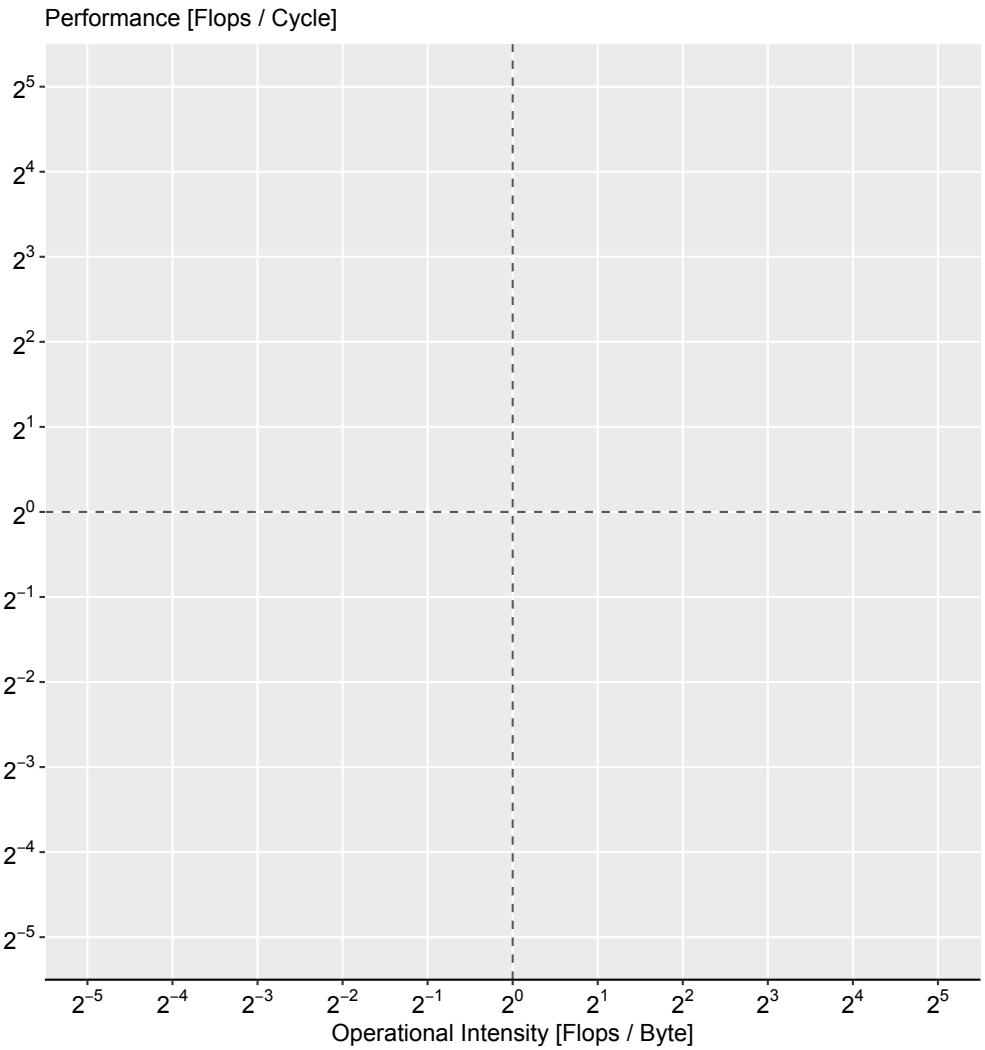

1. Draw the roofline plot for this computer into the above graph. Annotate the lines so we see your reasoning.

2. Consider the following computation where x, y, and z are arrays. Assume that x, y, and z are cache-aligned allocated (i.e., the address of an array maps with the first cache block):

```
1 void compute(float* x,float* y,float* z, int n){
2     for(int i=0; i < n; i++){
3         y[i+1] = (x[i]*y[i] + y[i]) + z[i];
4     }
5 }
```

   Based on the instruction mix (i.e. ignoring dependencies), which performance is maximally achievable for this function and why? Draw an associated tighter horizontal roofline into the plot above.

3. At what operational intensity $I(n)$ does this new horizontal roofline intersect with the memory roofline?

4. Based on the instruction mix **and** on data dependencies, which performance is maximally achievable for this function and why? We repeat the code for convenience:

```
1 void compute(float* x,float* y,float* z, int n){
2     for(int i=0; i < n; i++){
3         y[i+1] = (x[i]*y[i] + y[i]) + z[i];
4     }
5 }
```

5. Assume the cache is directly mapped, initially cold, and the expressions are evaluated from left to right ignoring data dependency. Assume the same array alignment as in Task 2.

   (a) Compute an upper bound (as tight as possible) for the operational intensity for large n. Consider only reads (i.e., ignore write-backs).

   (b) Based on this $I(n)$, which peak performance is achievable on the specified system taking into account instruction mix (i.e. the setting of Task 2)? In addition, indicate whether the computation is compute or memory bound.

6. What minimal associativity should the cache have to achieve the same performance as with a fully associative cache.

Problem 5: Cache Miss Analysis (18 = 7+5+6)

Consider the following computation that takes as input matrices X, Y and Z of size $n \times n$. `sizeof(double) = 8`. (X, Y, Z are not aliased).

```
1  /* NOTE: Assume that the notation A[i][j] is transformed to A[i*n + j].
2   *       We use the notation A[i][j] for readability only. */
3  void f(double *X, double *Y, double *Z, int n){
4    double t1, t2;
5    for (int i = 1; i < n-1; i += 1) {
6      for (int j = 1; j < n-1; j += 1){
7        t1 = X[i][j] + 0.25*(X[i][j-1] + X[i][j+1] + X[i-1][j] + X[i+1][j]);
8        t2 = Y[j][0] + Y[j][4] + Y[j][8];
9        Z[i][j] = t1 + t2;
10     }
11   }
12 }
```

Assume a fully associative write-back/write-allocate cache of size γ bytes, a cache block size of 64 bytes, and only one cache. Further assume that n is much larger than γ. In the following we perform two cache miss analyses assuming an initially cold cache. **In the derivations you can omit lower order terms** (writing $\approx$ instead of $=$). Show your work.

1. Considering cache misses from both reads and writes, estimate the number of cache misses incurred by `f` as a function of n.

2. Assume that we apply loop interchange, i.e., we swap the lines 5 and 6 in the code above. Estimate again the number of cache misses incurred by `f` when this optimization is applied. The code looks now as follows:

```
1  void f(double *X, double *Y, double *Z, int n){
2    double t1, t2;
3    for (int j = 1; j < n-1; j += 1){
4      for (int i = 1; i < n-1; i += 1) {
5        t1 = X[i][j] + 0.25*(X[i][j-1] + X[i][j+1] + X[i-1][j] + X[i+1][j]);
6        t2 = Y[j][0] + Y[j][4] + Y[j][8];
7        Z[i][j] = t1 + t2;
8      }
9    }
10 }
```

3. Now we try to reduce the number of misses by blocking the computation into blocks of size $b \times b$, b a multiple of 8. This means it now has the following loop structure (we ignore clean-up code if b does not divide $n-2$):

```
1  void f_blocked(double *X, double *Y, double *Z, int n){
2    double t1, t2;
3    for (int i = 1; i < n-1; i += b)
4      for (int j = 1; j < n-1; j += b)
5        for (int i1 = i; i1 < i+b; i1 += 1)
6          for (int j1 = j; j1 < j+b; j1 += 1){
7            t1 = X[i1][j1] + 0.25*(X[i1][j1-1] + X[i1][j1+1]
8                                   + X[i1-1][j1] + X[i1+1][j1]);
9            t2 = Y[j1][0] + Y[j1][4] + Y[j1][8];
10           Z[i1][j1] = t1 + t2;
11         }
12 }
```

Estimate the number of cache misses incurred by `f_blocked`. In doing so, upper bound the size of b so that you achieve good cache locality. Show also this bound.

Problem 6: Sampler (12 = 2+2+2+2+2+2)

Be brief in your answers, no need to show derivations.

1. You are given a computation of the form:

```
1 for(int i = 0; i < n; ++i) {
2   acc = acc OP x[i];
3 }
```

   where `OP` is a binary operation with gap = 0.5 cycles/issue and latency L. Assume that we unroll the loop by a factor of K and use K accumulators to improve ILP (the computation stays the same when $K = 1$). Assume that n is large. How many accumulators should you use at least to achieve the highest performance?

2. Give an example of a SIMD intrinsic that is not mapped to a single assembly instruction.

3. Explain what `_mm256_loadu_pd` does.

4. Consider the following computation where x and y are arrays.

```
void compute(float* x, float* y, int n){
  for(int i=0; i < n; i++){
    y[i] += x[i]*y[i];
  }
}

...
int reps = 100000;
uint64_t t1 = time();
for(int r=0; r < reps; r++){
  compute(x, y, N);
}
double mean_duration = (time() - t1)/(double)reps;
```

Assume $N = 128$. Would the code above provide a reliable measurement on existing modern CPUs, if we plan to measure the performance of the function for cold caches? Explain why. How can we modify the code to achieve that goal?

5. Consider a processor with a 16-way set associative cache of size 128KiB and block size 32 bytes. Propose a suitable page size (as small as possible) such that the cache look up can start concurrently with the TLB look up.

6. Provide col_idx(M) and row_start(M) of the below matrix when expressed in Compressed Sparse Row (CSR) format.

$$M = \begin{pmatrix} 0 & 1 & 2 & 0 \\ 0 & 3 & 0 & 0 \\ 9 & 0 & 5 & 0 \\ 7 & 0 & 0 & 4 \end{pmatrix}$$